

How to Write Good Reviews for CVPR

by CVPR 2019 Program Chairs
Derek Hoiem, Gang Hua, Abhinav Gupta, and Zhuowen Tu

Outline – what you can learn

- Why your job as reviewer is so important
- How the paper decision process works
- How to structure a review, with good and bad examples
- Tips for reviewers

Thank you for serving as a reviewer!

We are all counting on you.

- Area chairs for **clearly justified guidance for paper accept/reject decisions**
- Authors for **fair consideration** and constructive feedback
- Community for ensuring that **every CVPR paper teaches something worthwhile**

If you write bad, poorly justified, ill-considered, or unfair reviews.....

- Your area chairs, who may greatly influence your career advancement, may remember that you let them down
- Authors may feel unwelcome or mistreated by the review process
- A reader may waste time on a flawed or uninformative paper that was accepted, or may waste time in research because a valuable paper was rejected

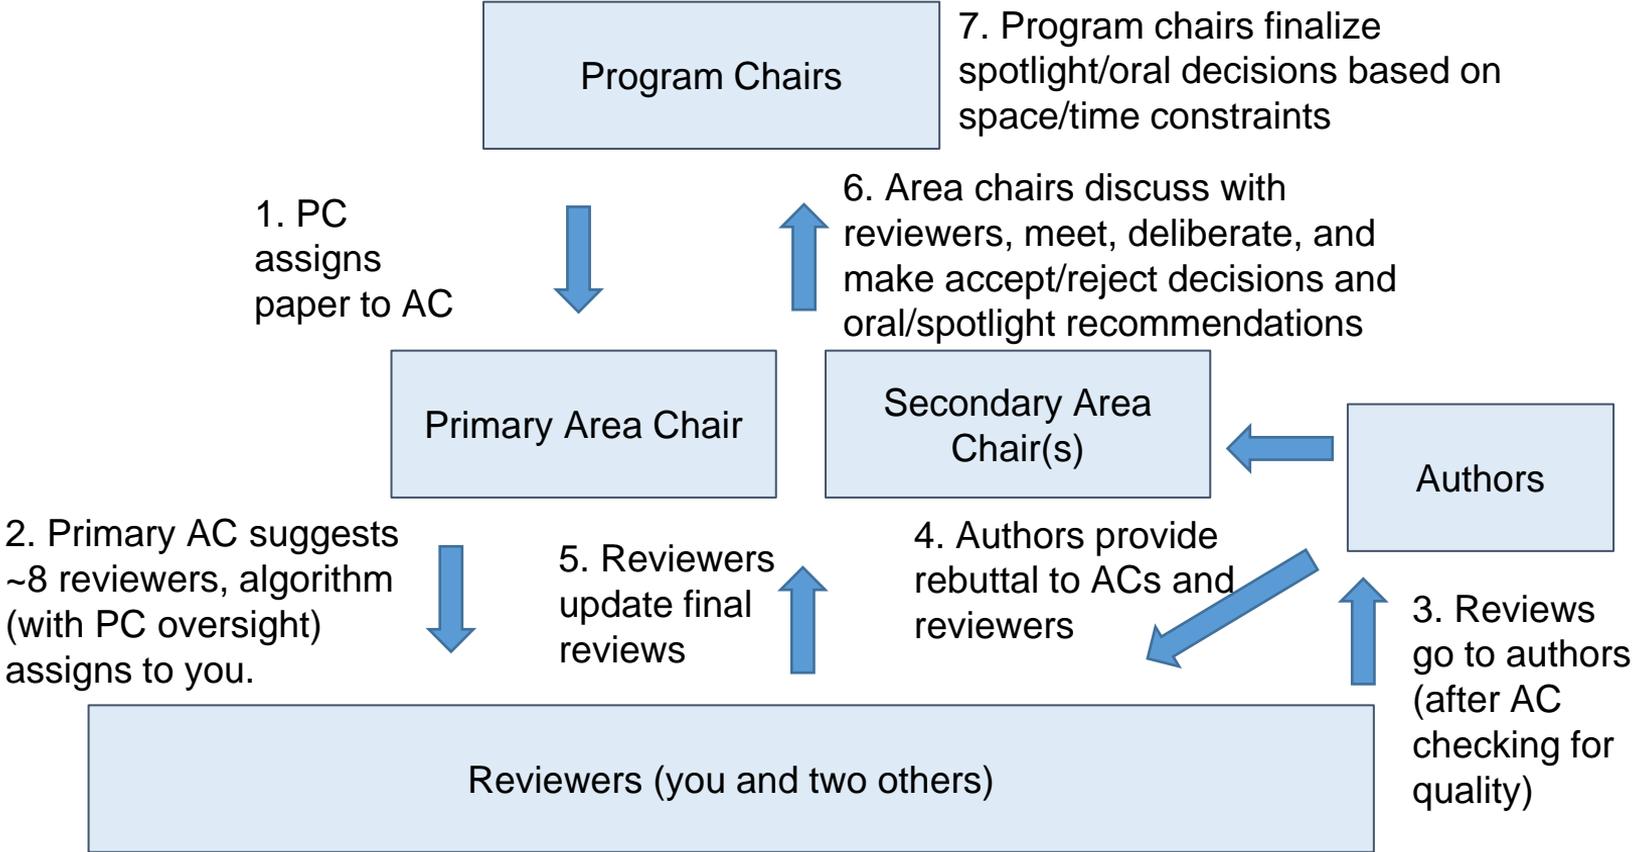
If you write good, insightful, well-justified, constructive reviews....

- Your ACs will love you because you will make the paper decision much easier
- The authors' faith in the vision community will increase, and, even if they need to resubmit, they will know what needs to improve
- The community will continue to flock to CVPR for the latest and greatest in computer vision ideas and techniques

So how do you write good reviews?

- Understand the process
- Learn the structure of a good review
- Put in the considerable effort required to understand each paper, the relevant literature, and to write a clear review

The decision process (overview)



The decision process (in detail)

1. **Program chairs (PCs) assign papers to area chairs (ACs)**, typically ~35 papers per area chair, 3300 total at CVPR'18
2. **ACs nominate ~8 reviewers per paper**, with some help from software
3. **Papers are assigned to reviewers** (3 per paper) using an optimization algorithm that takes into account area chair suggestions, paper load constraints, conflict constraints, and the prayers of program chairs that nothing goes wrong.
4. **Reviewers submit initial reviews**, typically handling 6-12 papers each
5. **ACs and PCs** check that reviews are complete and not obviously negligent
6. **Authors receive reviews** with a mixture of gasps, grimaces, grumbles, and the occasional grin. After much thought and re-reading of paper and reviews, authors submit rebuttals.
7. **Discussion ensues among reviewers and AC**, based on all reviews, rebuttal, and paper. Reviewers update their ratings and justification.
8. **ACs make initial decisions** (usually the easy papers, where all the reviewers agree on accept or reject and justification is clear)
9. **All accept/reject decisions are finalized at the area chair meeting.** ACs form buddies or triplets, within larger panels. There are no conflicts within a panel. The decision and summary are recorded by the primary AC for each paper and checked/approved by a secondary AC. Papers with difficult decisions may be read by others in the triplet or panel, and additional opinions may be sought by other ACs after checking for conflicts. In addition to accept/reject decisions, panels provide a roughly ranked list of oral/spotlight nominations to the PCs.
10. **PCs make final determination of poster/spotlight/oral** for accepted papers, almost entirely based on the recommendations of the ACs but taking into account time and space constraints and topic diversity.

How to structure a review

Summary: Explain the key ideas, contributions, and their significance. This is your abstract of the paper. The summary helps the AC understand the rest of your review and be confident that *you* understand the paper.

Strengths: What about the paper provides value -- interesting ideas that are experimentally validated, an insightful organization of related work, new tools, impressive results. Most importantly, what can someone interested in the topic learn from the paper.

Weaknesses: What detracts from the contributions -- Does the paper lack controlled experiments to validate the contributions? Are there misleading claims or technical errors? Is it possible to understand (and ideally reproduce) the method and experimental setups by reading the paper?

Rating and Justification: Carefully explain why the paper should be accepted or not. This section should make clear which of the strengths and weaknesses you consider most significant.

Examples

- The following examples are from ICLR, which published reviews in the public domain
- For ICLR, the review is written as a single statement, rather than broken into sections as for CVPR, but the same criteria for reviews apply
- Here we consider the **quality of the form**, rather than the accuracy of the content, of the review. Of course the accuracy of the review is also very important but requires much more expertise and time to analyze

Review quality: Good. Though missing a summary of contribution, the review clearly explains why the paper should be accepted

(Note: this was a late-added review, which may account for brevity)

Rating: 9: Top 15% of accepted papers, strong accept

Review: First off, this paper was a delight to read. The authors develop an (actually) novel scheme for representing spherical data from the ground up, and test it on three wildly different empirical tasks: Spherical MNIST, 3D-object recognition, and atomization energies from molecular geometries. They achieve near state-of-the-art performance against other special-purpose networks that aren't nearly as general as their new framework. The paper was also exceptionally clear and well written.

The only con (which is more a suggestion than anything)--it would be nice if the authors compared the training time/# of parameters of their model versus the closest competitors for the latter two empirical examples. This can sometimes be an apples-to-oranges comparison, but it's nice to fully contextualize the comparative advantage of this new scheme over others. That is, does it perform as well and train just as fast? Does it need fewer parameters? etc.

I strongly endorse acceptance.



+ Clearly explains why the paper should be accepted
– Does not contain many details about the contribution or why it is novel, so relies on the AC trusting the reviewer's judgment on these points

Note: though the proposed method does not achieve the best results (according to the review), the paper is highly valued for proposing a more general framework. Achieving best results is not necessary to validate the key idea (e.g., generality by testing with diverse datasets, or including an ablation study that isolates the impact of the key idea).

+ Indicates that the reviewer tried to think of weaknesses but could not come up with anything that should negatively impact the paper rating
+ Constructive feedback for the authors

Review Quality: OK but not great. Makes general factors in decision clear and provides detailed feedback to authors, but does not provide adequate explanation for strengths and weaknesses

Rating: 8: Top 50% of accepted papers, clear accept

The paper proposes a framework for constructing spherical convolutional networks (ConvNets) based on a novel synthesis of several existing concepts. The goal is to detect patterns in spherical signals irrespective of how they are rotated on the sphere. The key is to make the convolutional architecture rotation equivariant.

Pros:

- + novel/original proposal justified both theoretically and empirically
- + well written, easy to follow
- + limited evaluation on a classification and regression task is suggestive of the proposed approach's potential
- + efficient implementation

Cons:

- related work, in particular the first paragraph, should compare and contrast with the closest extant work rather than merely list them
- evaluation is limited; granted this is the nature of the target domain

Presentation:

- * While the paper is generally written well, the paper appears to conflate the definition of the convolutional and correlation operators? This point should be clarified in a revised manuscript.
- * In Section 5 (Experiments), there are several references to S^2 CNN. This naming of the proposed approach should be made clear earlier in the manuscript. As an aside, this appears a little confusing since convolution is performed first on S^2 and then $SO(3)$.

Evaluation:

- * What are the timings of the forward/backward pass and space considerations for the Spherical ConvNets presented in the evaluation section? Please provide specific numbers for the various tasks presented.
- * How many layers (parameters) are used in the baselines in Table 2? If indeed there are much less parameters used in the proposed approach, this would strengthen the argument for the approach. On the other hand, was there an attempt to add additional layers to the proposed approach for the shape recognition experiment in Sec. 5.3 to improve performance?

Minor Points:

- some references are missing their source, e.g., Maslen 1998 and Kostolec, Rockmore, 2007, and Ravanbakhsh, et al. 2016.
- [abridged minor points due to lack of space in this slide]
- Figure 5, caption: "The red dot correspond to" -> "The red dot corresponds to"

Final remarks:

Based on the novelty of the approach, and the sufficient evaluation, I recommend the paper be accepted.



+ Highlights key ideas and contributions.

- The summary should also include one sentence on experimental setup

- Summary should include one sentence on significance of the contribution



+ Itemizes strengths and weaknesses

- Does not provide enough detail. E.g., what is original about the paper? How is the evaluation limited?



+ Includes clarifications questions and constructive feedback for authors

+ Makes it clear that "Minor Points" are not an important factor in decision



+ Identifies key positive factors in rating

- Would have been better to say why the weaknesses are given less weight

Review quality: Bad. The review lists only weaknesses and requests for clarification, omitting a summary and justification for decision. Thus, it is unclear to author or AC which of these points are the primary basis for the rating.

Rating: 4: Ok but not good enough - rejection

Review: 1. The idea of multi-level binarization is not new. The author may have a check at Section "Multiple binarizations" in [a] and Section 3.1 in [b]. The author should also have a discussion on these works.

2. For the second contribution, the authors claim "Temperature Adjustment" significantly improves the convergence speed. This argument is not well supported by the experiments.

I prefer to see two plots: one for Binarynet and one for the proposed method. In these plot, testing accuracy v.s. the number of epoch (or time) should be shown. The total number of epochs in Table 2 does not tell anything.

3. Confusing in Table 2. In ResBinNet, why 1-, 2- and 3- level have the same size? Should more bits required by using higher level?

4. While the performance of the 1-bit system is not good, we can get very good results with 2 bits [a, c]. So, please also include [c] in the experimental comparison.

5. The proposed method can be trained end-to-end. However, a comparison with [b], which is a post-processing method, is still needed (see Question 1).

6. Could the authors also validate their proposed method on ImageNet? It is better to include GoogleNet and ResNet as well.

7. Could the authors make tables and figures in the experiment section large? It is hard to read in current size.

Reference

[a] How to Train a Compact Binary Neural Network with High Accuracy. AAAI 2017

[b] Network Sketching: Exploiting Binary Structure in Deep CNNs. CVPR 2017

[c] Trained Ternary Quantization. ICLR 2017

+ Cites papers that make the idea "not new"

– Does not say how these methods relate, so it is not clear if they are very similar techniques

– Because it is not tested by experiments, or that the convergence speed is not different?

- The remaining points may help authors improve the paper, but it is not clear if they are a significant factor in the rating to reject

Big problems:

- AC can't make good use of the review without reading the paper, due to lack of summary/justification.
- No strengths listed, which may indicate that reviewer is just looking for reasons to reject.
- Author and AC don't know which of the listed points are important for reject rating.

Common issues

How to treat supplemental material

The supplemental material is intended to provide details of derivations and results that does not fit within the paper format or space limit. It is not an extension of the deadline. The paper should indicate which materials are in the supplemental material, and you need consult only if you think it is helpful in understanding the paper and its contribution.

How to handle policy violations (anonymity, plagiarism, scope, etc.)?

- (1) Contact the program chairs with the paper number and explanation of the suspected problem.
 - (2) Review the paper as if there is no problem
- Program chairs will follow up on the issue, but it may take some time.

No Additional Experiments in Rebuttals

Per a 2018 PAMI-TC motion, reviewers should not request additional experiments for the rebuttal, or penalize for lack of additional experiments. Authors should not include new experimental results in the rebuttal.

Tips for reviewers

- **The ACs are your primary audience.** Make your review **self-contained**, and clearly justify your opinions and ratings. A good review describes the key ideas, strengths, and flaws in a way that can be understood by someone who has not yet read the paper.
- **Take the time to do a good review.** Many experienced reviewers take 2-4 hours per paper. If you're fairly new to reviewing (e.g. grad student), plan on at least 4 hours per paper and take the time to read the paper twice, consider related work, look up unfamiliar techniques, etc
- **Be conscious of your own bias.** We all tend to assign more value to papers that are relevant to our own research. Try to ignore “interestingness of topic” and focus on whether the paper can teach something new to an interested reader